Paper AAS 02-121

Design of the Reaction Wheel Attitude Control System for the Cassini Spacecraft

Glenn A. Macala

Jet Propulsion Laboratory
California Institute of Technology
M.S. 198-326
4800 Oak Grove Drive
Pasadena, CA 91109

AAS/AIAA Space Flight Mechanics Meeting

San Antonio, Texas **27-30 January, 2002**

AAS Publications Office, P.O. Box 28130, San Diego, CA 92129

Design of the Reaction Wheel Attitude Control System for the Cassini Spacecraft*

Glenn A. Macala, Jet Propulsion Laboratory

This paper presents the architecture and methodology used for the design of the Cassini spacecraft's Reaction Wheel Attitude Controller. Simulation results are presented that predict the pointing performance. Preliminary in-flight pointing performance results are also presented.

INTRODUCTION

The Cassini spacecraft (S/C) was launched in October 1997 on a mission to Saturn. In July 2004 the S/C will achieve orbit about Saturn where it will use its remote-sensing instruments to study the Saturnian system for 4 years. In addition, it will deliver an atmospheric probe (Huygens) into the atmosphere of Saturn's largest moon, Titan, and relay the resulting data from the probe back to Earth.

Cassini uses two types of attitude control systems at different portions of its mission (see Figure 1). A set of reaction control thrusters (RCS) is used for S/C detumble and momentum management, and for attitude control during lengthy cruise periods and during S/C safing. A reaction wheel assembly (RWA) is used primarily for attitude control when precision pointing of the remote science instruments is required during the nominal 4 year mission at Saturn and also when gravitational wave experiments are carried out during certain portions of the cruise to Saturn.

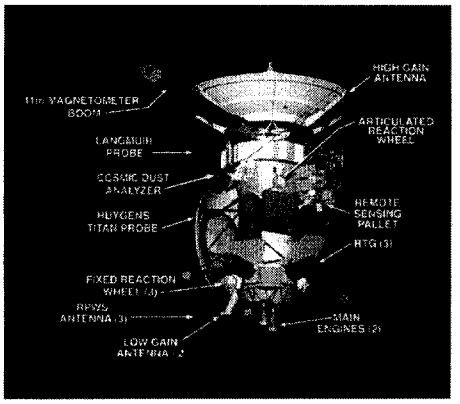

Figure 1. Cassini Spacecraft

In RWA Attitude Control mode, the S/C uses an algorithm called the Reaction Wheel Attitude Controller (RWAC) to control S/C attitude and rate. The RWAC must be able to keep S/C attitude inertially fixed for certain science observations. Additionally, the RWAC must be able to control S/C attitude and rate during S/C turns and target motion compensation, in which case the S/C inertial attitude and attitude rate may be required to vary with time.

This paper describes the design of the Reaction Wheel Attitude Controller. In addition, preliminary performance results from an early use of the RWAC at the Jupiter flyby are presented.

REQUIREMENTS

The pointing requirements of concern to the RWAC relate primarily to pointing the boresight of the Remote Sensing Pallet, but they are specified relative to the S/C coordinate axes. Referring to Figure 1, the S/C Z axis is anti-parallel to the High Gain Antenna boresight, the S/C Y axis is parallel to the Magnetometer Boom, and the S/C X axis completes the right-handed triad. The boresight of the Remote Sensing Pallet is aligned with the negative S/C Y axis (anti-parallel to the S/C Magnetometer Boom).

The primary pointing accuracy requirements levied on the RWAC can be summarized as a combination of two items:

1. Relative pointing errors due solely to transients and operating characteristics of the RWAC are allowed to contribute as much as 40 μrad each to the absolute S/C attitude error in X, Y, and Z axes at maximum slew rates of 0.13, 0.13, and 0.26 deg/sec in X, Y, Z axes, respectively.

2. Absolute S/C pointing stability at selected time windows shall be per Table 1, "S/C RMS Per-Axis Pointing Stability Requirements". Pointing stability requirements and capabilities reported herein are applicable only during instrument staring (i.e., zero inertial slew rate).

Table 1
S/C RMS Per-Axis Pointing Stability[1] Requirements

Time Window (sec)	2 σ per axis (μrad)
0.5	4
1	8
5	36
22	100
100	160
900	200
1200	220
1 hr	280
4 hr	800

As noted in Requirement 2, the requirements of Table 1 shall apply only at zero inertial slew rates. An allowable settling time after the completion of a S/C rate change shall be less than 20 seconds when the S/C accelerating torque is below 0.03 Nm.

DESIGN

Because the S/C X, Y, and Z axes lie approximately along the principal axes of the S/C, the basic architecture for the RWAC is a decoupled 3-axis, Proportional-Derivative (PD) discrete-time controller executing at 8 Hz controlling the orientation of a rigid body. The decoupling is further enforced by

commanding control torques using Euler's equation for the dynamics of a rigid body plus a set of reaction wheels as shown in the following vector equation.

$$Torque_{cmd} = I_{S/C}\dot{\omega}_{cmd} + \omega_{S/C}x(I_{S/C}\omega_{S/C} + H_{whls})$$

If the S/C inertia tensor, $I_{S/C}$, were exactly known along with the S/C's angular rate vector, $\omega_{S/C}$, and the momentum vector of the reaction wheels, H_{whls}, then application of the torque indicated by the above equation would result in a S/C angular acceleration exactly equal to that commanded by the RWAC, $\dot{\omega}_{cmd}$.

Figure 2 is an illustration of this decoupling concept. Note that the far right in the diagram shows what appears to be a single axis plant reacting to a torque command and disturbance inputs. This plant is actually an implementation of the above equation, but is shown in the block diagram as single axis to more simply illustrate the basic control loop concepts.

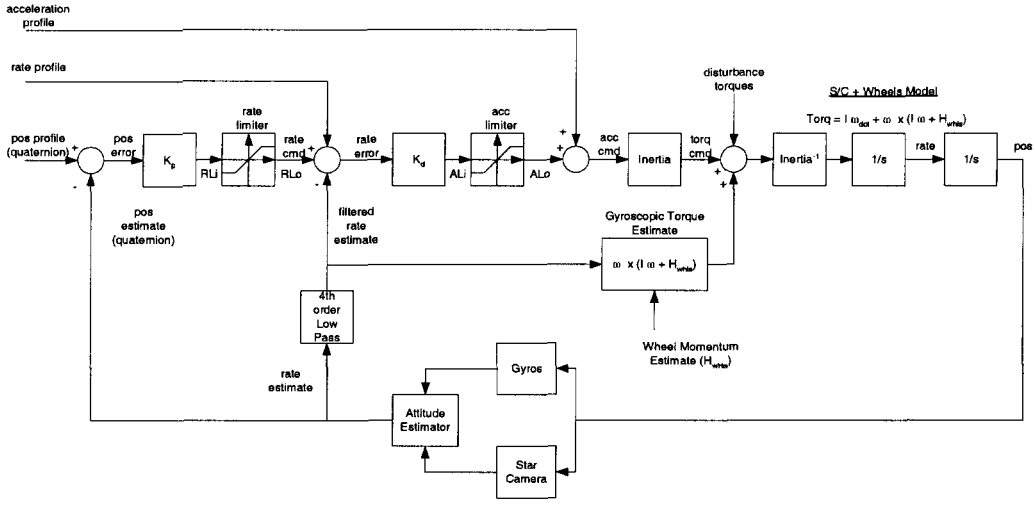

Figure 2 Block Diagram of the RWAC

The next basic control loop concept used in the RWAC is to provide feedforward angular acceleration and rate commands along with the attitude command. The feedforward acceleration command is multiplied by the S/C's inertia tensor in order to compute the required accelerating torque command. This command is added to the gyroscopic torque compensation and then passed along to the reaction wheels for application to the S/C. If every component of this chain were ideal, the S/C would achieve the desired acceleration. This acceleration would result in an angular rate equal to the commanded rate and an attitude equal to the commanded attitude. In other words, the attitude and rate profiles would be achieved with no further control action required.

This concept is further enforced by making sure that the attitude, rate, and acceleration commands are all consistent and physically realizable by the S/C. The construction of these quantities is handled by another algorithm onboard the S/C called the Attitude Commander[2]. Simple constant acceleration/deceleration or acceleration/coast/deceleration attitude profiles are computed in response to attitude commands sent to the S/C by mission operations. Furthermore, these profiles are passed through an algorithm called the Constraint Monitor[3] before being sent to the RWAC. Among other things, the Constraint Monitor checks the profiles using acceleration limits and rate limits that are compatible with reaction wheel torque

capability. If a violation is found, the Constraint Monitor substitutes compatible profiles with the hope that the violation is temporary and attitude will eventually catch up with the commanded attitude.

At this point in the design, the attitude control has been constructed to work perfectly in the absence of attitude feedback. However, we know that this is an idealization that can never be achieved. We therefore add feedback from an Attitude Estimator. The Attitude Estimator uses gyros and a star camera in a Kalman filter to estimate inertial attitude and attitude rates.

The attitude is computed as a quaternion. A quaternion multiplication with the commanded quaternion (attitude) is performed to find the attitude error. This error is then multiplied by a gain, K_p, to form a rate command that will remove this attitude error. The rate command is limited (more about this later) and added to the rate profile. This total rate command is then compared to the estimated rate and the error is also multiplied by a gain, K_d, to form an acceleration command that will remove this rate error. The acceleration command is then limited (more about this later) and added with the acceleration profile to form the total acceleration command. As discussed earlier, this acceleration command is multiplied by the S/C inertia tensor to compute the wheel torques to be applied to the S/C.

In the absence of any significant bias torques acting on the S/C, this PD controller should remove all attitude errors. However, because of bearing friction, the reaction wheels have drag torques acting upon them. The spinning wheels will require a compensating motor torque to be applied in order to keep them spinning at a constant rate and to exactly cancel the resulting reaction torques being applied to the S/C. If the RWAC were to provide this torque, a constant attitude error would result due to the PD structure of the controller. Attitude error would be required to be equal to bearing friction divided by the quantity $K_p*K_d*I_{S/C}$. This is clearly undesirable.

We could have increased the complexity of the RWAC by adding integral control to the design. However, because the only significant drag torques acting upon the S/C are the reaction wheel bearing torques, we decided to place the compensation for the torques in the Reaction Wheel Hardware Manager algorithm. This algorithm performs other tasks related to wheel control such as wheel speed estimation using wheel tachometers. We therefore put a local control loop around each wheel that uses a Proportional/Integral (PI) controller to drive wheel speed to commanded wheel speed. The controller passes wheel torque commands from the RWAC straight through to the wheel motor for application. These commanded torques are also used to compute the ideal wheel speeds that would result from their application. Estimated wheel speeds are computed by low pass filtering the wheel tachometers. The ideal wheel speeds are then compared to the estimated wheel speeds and the errors are sent to the PI wheel controller. The PI controller augments the RWAC wheel torque commands with additional torque commands that compensate for bearing torques and other wheel imperfections. This scheme effectively places integral control locally about each wheel, removing the need for integral control in the RWAC.

With the basic structure of the RWAC in place, we can now determine its required bandwidth. A principal driver for bandwidth is the requirement that "An allowable settling time after completion of a S/C rate change shall be less than 20 seconds ...". This puts a requirement on the time constant of the RWAC. If we model the response of the RWAC to a transient attitude disturbance (θ_d) as a decaying exponential with a time constant equal to the inverse of the RWAC's bandwidth, then the required bandwidth would be approximately equal to 4 times this time constant. In equation form,

$$AttitudeTransient \approx \theta_d e^{-\omega_n t}$$

Then to remove 98% (4 time constants) of the disturbance in attitude in about 20 seconds requires that the bandwidth, ω_n, be approximately:

$$\omega_n \approx 4/20 rad/\sec \approx 0.03 Hz$$

With this bandwidth in mind, we can now select control loop gains and evaluate the effects of propellant slosh and structural flexibility on the design.

ANALYSIS

If we approximate the RWAC as a one-axis, linear, continuous-time version of the system shown in Figure 2, the 2nd order characteristic polynomial of this system can be used to solve for its roots in terms of gains K_p and K_d. The result is written below using the damping factor, ζ, and natural frequency, ω_n, of the roots (poles) of a 2nd order system:

$$K_p = \omega_n/(2\zeta)$$

$$K_d = 2\zeta\omega_n$$

At this point we could substitute a value for ω_n (2π 0.03) and ζ (~0.707) and evaluate our design in terms of gain and phase margins and pointing results. This was done using a rigid body model and the results were acceptable. However, before we can proceed with this design, we must also consider the effects of propellant slosh and structural flexibility on our design.

The S/C has 2 large propellant tanks that feed the main engine oxidizer and fuel. It turns out that the propellant in these tanks can be modeled as pendulums constrained by spring/damper sets while the RWAC is controlling the S/C. This is mainly due to the action of propellant management devices in the tanks. The natural frequencies of these pendulums are around 0.004 Hz or so. Since these frequencies are well within our bandwidth of 0.03 Hz and are appendage-like modes (zero/pole pairs), these modes should not pose a problem for the RWAC.

The S/C also has a long Magnetometer Boom with a natural frequency around 0.65 Hz. This mode has a rather large structural gain. Therefore, even though this frequency is well outside of the RWAC's bandwidth, a low pass filter is needed to attenuate its gain in the feedback loop. Since the mode's structural gain is about 35 dB, we would like at least that much attenuation from the filter at that frequency. Because most of the high frequency gain comes back through the loop via the rate feedback, we placed the low pass filter into the rate feedback portion of the RWAC.

Lastly, the S/C has a set of light-weight but long (~10 m) radio frequency booms (RPWS). The RPWS booms have a natural frequency at about 0.13 Hz. This frequency lies much closer to our bandwidth and is in a region where the phase stabilization of the mode is important. Furthermore, there is uncertainty in the modal frequency. We therefore require that the low pass filter does not destabilize this mode.

Given the discussion above, the approach taken for final gain selection and filter selection can be summarized as follows:

1. Construct the frequency response using a PD control law with a full cycle computational delay controlling a rigid body plant. Assume that the gyro contributes little at frequencies less than 1 Hz, so ignore it in the design. However, include a low-pass rate filter: first try a 2nd order filter, and if more gain reduction is needed at the magnetometer boom frequency, try a 4th order filter.

2. Since the rigid body design was performing adequately in terms of pointing performance and magnitude of control torques, hold the gain at low frequency, say at 0.003 Hz (low-g slosh area), at its current value of 40 dB.

3. Constrain the gain at 0.65 Hz to be less than -47 dB to allow for at least 10 dB of gain margin for the magnetometer boom mode.

4. Try to get at least 20° of phase margin at 0.2 Hz for the RPWS mode.

5. Limit the natural damping (ζ) in the PD design and in the low pass filter to no less than 0.4 to reduce the effects of transients in the control loop.

6. Limit the control loop's rigid body phase margin to no less than 40°, again to reduce the effects of transients in the control loop.

7. Tune the natural frequency and damping of the PD loop and the natural frequency and damping of the low-pass filter to meet objectives 2 through 6.

Using the NPSOL[4] program, the above constrained optimization problem was solved. The results were such so as to require that a fourth order rate filter be used in order to achieve the objectives. The resulting, "tuned" design is illustrated in Figure 3.

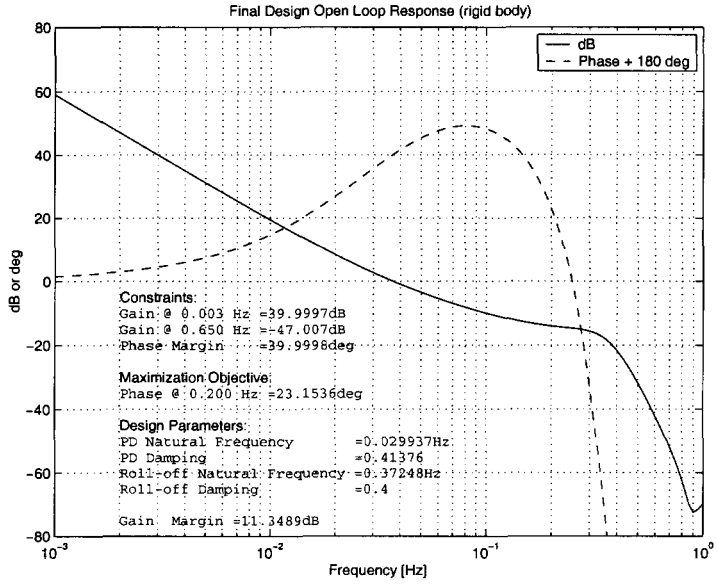

Figure 3 RWAC Open Loop Frequency Response

Since the above design procedure utilized a rigid body plant model to simplify the optimization procedure, we must now confirm the results using the flexible body plant: the S/C model including propellant slosh, RPWS flexible booms, and Magnetometer flexible boom models. Figure 4 shows the results of combining the tuned design with the flexible body model (the X-axis RWAC loop in this case). We see that we have met our objectives: > 40° rigid body phase margin; > 12 dB rigid body gain margin; at least 20° phase margin at 0.2 Hz to protect RPWS stability; > 10 dB gain margin at the Magnetometer Boom resonance. Naturally these margins were verified for X, Y, and Z RWAC channels, including expected variations in the models.

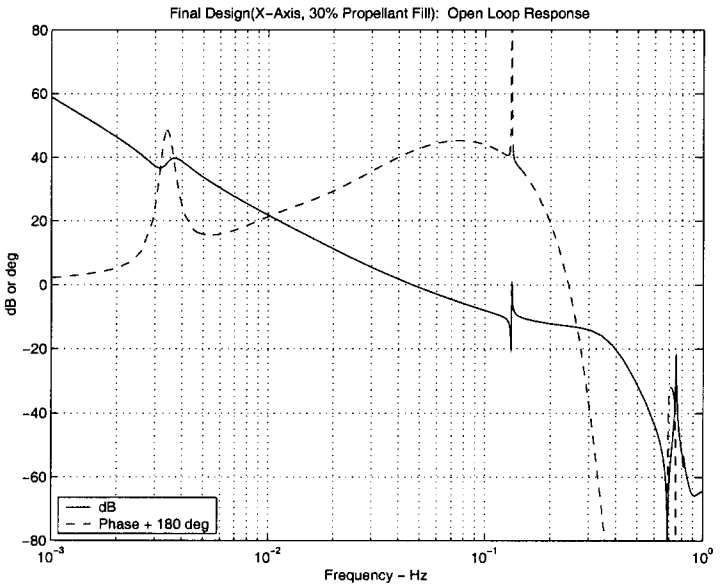

Figure 4 RWAC Open Loop Frequency Response: Full S/C Model

DYNAMIC LIMITERS

One last constraint was imposed on the design of the RWAC. Because the reaction wheels have a very small torque capability (~0.16 Nm) with respect to the S/C inertia tensor (~8000 kg-m^2), the angular acceleration capability of the S/C is quite small. Furthermore, not all of the reaction wheel torque is available for accelerating the S/C. The total wheel torque must be budgeted to cover S/C accelerating torque, wheel bearing drag torque, gyroscopic torque, and attitude control torques used to remove attitude control errors.

The amount of torque available to the RWAC to remove attitude control errors was budgeted to be about 0.02 Nm. To enforce this budget, an acceleration limiter was inserted into the RWAC (*in Figure 2, ALi marks its input, ALo marks its output*). The acceleration limit is set such that its product with the S/C inertia results in no more than 0.02 Nm of commanded torque to the wheels. Because acceleration is limited, large overshoots in attitude could occur if rate were not limited also. To prevent this occurrence, a rate limiter was inserted into the RWAC (*in Figure 2, RLi marks its input, RLo marks its output*). The rate limiter was set to the value of the acceleration limit divided by K_d. In other words, the rate limiter was set to a value that would just cause the acceleration limiter to be hit.

This arrangement performed well even though the rate limit was very small (~ 20 μrad/sec). However, testing of the flight software in a real-time testbed showed that if a large attitude error somehow developed, removing that error at 20 μrad/sec could take very large amounts of time. It turned out that the real-time testbed contained imperfections that occasionally caused large attitude errors to be developed. Since the flight software would be operating on the real S/C, not the imperfect real-time testbed, this behavior should not be seen in flight. However, it was decided that if ever a large attitude error developed *and* if any portion of the S/C accelerating torque budget and S/C rate budget were not currently being used, the RWAC should be permitted to make use of them to more quickly reduce those errors.

Therefore, the acceleration limiter of Figure 2 (with input ALi and output ALo) was replaced with the adjustable acceleration limiter shown in Figure 5. The adjustable limiter can vary its limit values each

compute cycle. It has a small value, accLowLimit, that is set to the acceleration value corresponding to the originally budgeted 0.02 Nm torque. In addition, each compute cycle the RWAC monitors the acceleration profile coming from the Constraint Monitor. If the magnitude of the current value of the acceleration profile is less than the maximum amount budgeted for S/C acceleration, accProfileMax, the excess is made available for use by the RWAC. This is implemented by setting the RWAC's acceleration limit to the sum of the small limit and the excess available from the acceleration profile every compute cycle.

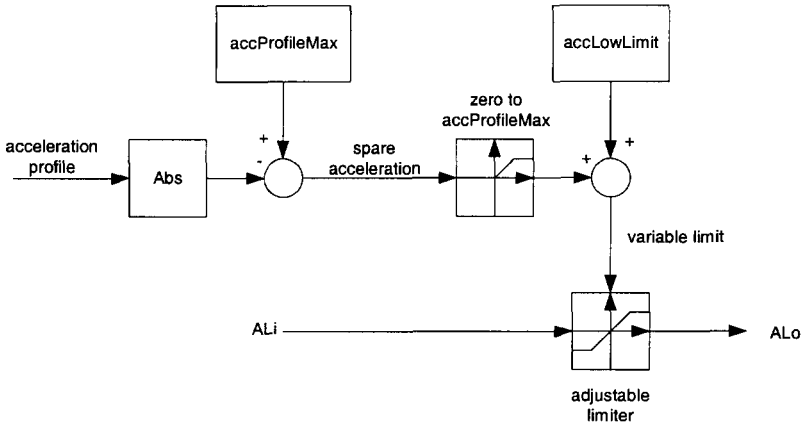

Figure 5 Adjustable Acceleration Limiter

The rate limiter of Figure 2 (with input RLi and output RLo) was replaced with the adjustable rate limiter shown in Figure 6. As with the adjustable acceleration limiter, the adjustable rate limiter can vary its limit values each compute cycle. It has a small value, rateLowLimit, that is set equal to the rate that would cause the smallest acceleration limit to be hit in the adjustable acceleration limiter. In addition, at each compute cycle the RWAC monitors the rate profile coming from the Constraint Monitor. If the magnitude of current value of the rate profile is less than the maximum amount budgeted for S/C rate, rateProfileMax, the excess is made available for use by the RWAC. This is implemented by setting the RWAC's *available* rate limit to the sum of the small limit and the excess available from the rate profile every compute cycle. However, another condition is necessary before the *actual* rate limit is allowed to be set higher.

We only wish to open up the rate limit if a large attitude error is present and has been persistent. To accomplish this, we pass a signal proportional to the square root of the magnitude of the raw rate feedback command (at RLi in Figure 2) through a low pass filter. The output of the filter is limited to rateProfileMax. The constant of proportionality is set to a value that would result in a commanded rate limit equal to a value such that given the current attitude error, the minimum acceleration (accLowLimit) would be enough to reduce a S/C rate from the value of the commanded rate limit to zero without an attitude overshoot. That value can be computed by assuming constant deceleration, single-axis rigid body motion. It turns out to be:

$$G_o = \sqrt{2\ accLowLimit / K_p}$$

The values for the digital low pass filter are set to achieve a DC gain of 1 and a bandpass frequency of about 0.05 Hz since we want the rate limit to be insensitive to attitude noise. The value of 0.05 Hz allows for a filter time constant of about 3 seconds. In other words, a significant attitude error would have to persist for at least 3 seconds before the rate limit could begin to open up. We also subtract a small bias from the persistent attitude error to account for any attitude noise being rectified by the absolute value, square root operation.

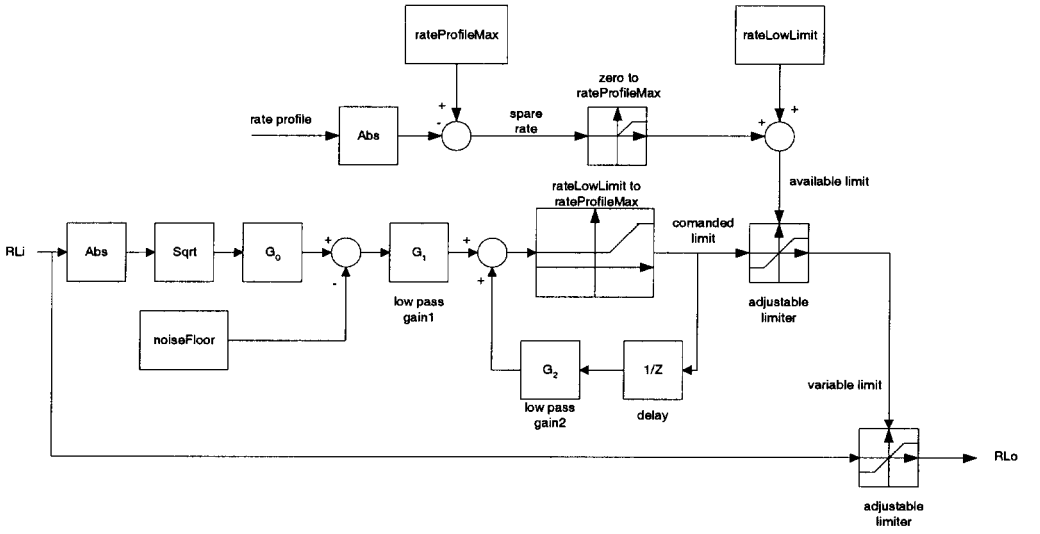

Figure 6 Adjustable Rate Limiter

Although these new dynamic limiters address the problem of removing large, anomalous attitude errors quickly, they add complexity to the RWAC and make it very difficult to rigorously analyze for stability. Because of this, a large number of time simulations were run to investigate the behavior of the RWAC under a variety of anomalous conditions. The results of one of those simulations are shown in Figure 7. This case is particularly interesting because it exercises most of the features of the dynamic limiters.

The top two graphs show the commanded and achieved S/C X-axis attitude and rate for a commanded constant rate turn about the X-axis. The S/C X-axis rate was intentionally set to an anomalous initial rate of 0.875 mrad/sec so as to trigger action in the dynamic limiters. Because the commanded slew profiles assume the initial S/C rate to be zero, both attitude and rate errors develop as shown in the middle two graphs. The bottom two graphs show the response of the dynamic limiters. Noted at the top of these graphs are the minimum and maximum allowable values for the dynamic limits. As position error persistently grows, the rate limiter exponentially opens up (bottom left, 0 to ~180 secs). During this same time, acceleration (bottom right) is limited to a little more than its minimum because the acceleration profile during this time is using a little less than its budget allows. This frees up some spare acceleration for the RWAC. As 200 seconds approach, we see the dynamic rate limit reducing. This is because the rate profile has reached values that now deplete the excess rate available to the dynamic rate limiter. At about this time, the rate profile assumes a constant rate and therefore the acceleration profile goes to zero, freeing up excess acceleration for the dynamic acceleration limiter to use. Note that the acceleration limit goes to its maximum value at that time.

From about 210 seconds until about 600 seconds, the constant rate profile leaves only a small amount of excess rate for the dynamic rate limiter. As attitude error is reduced (at the limited rate), eventually it becomes small enough to cause the dynamic rate limiter to reduce its limit to its lower limit value. Note that attitude error suffers no overshoot, as the design intended. At this time, position errors are essentially zero and the anomalous initial condition has been removed.

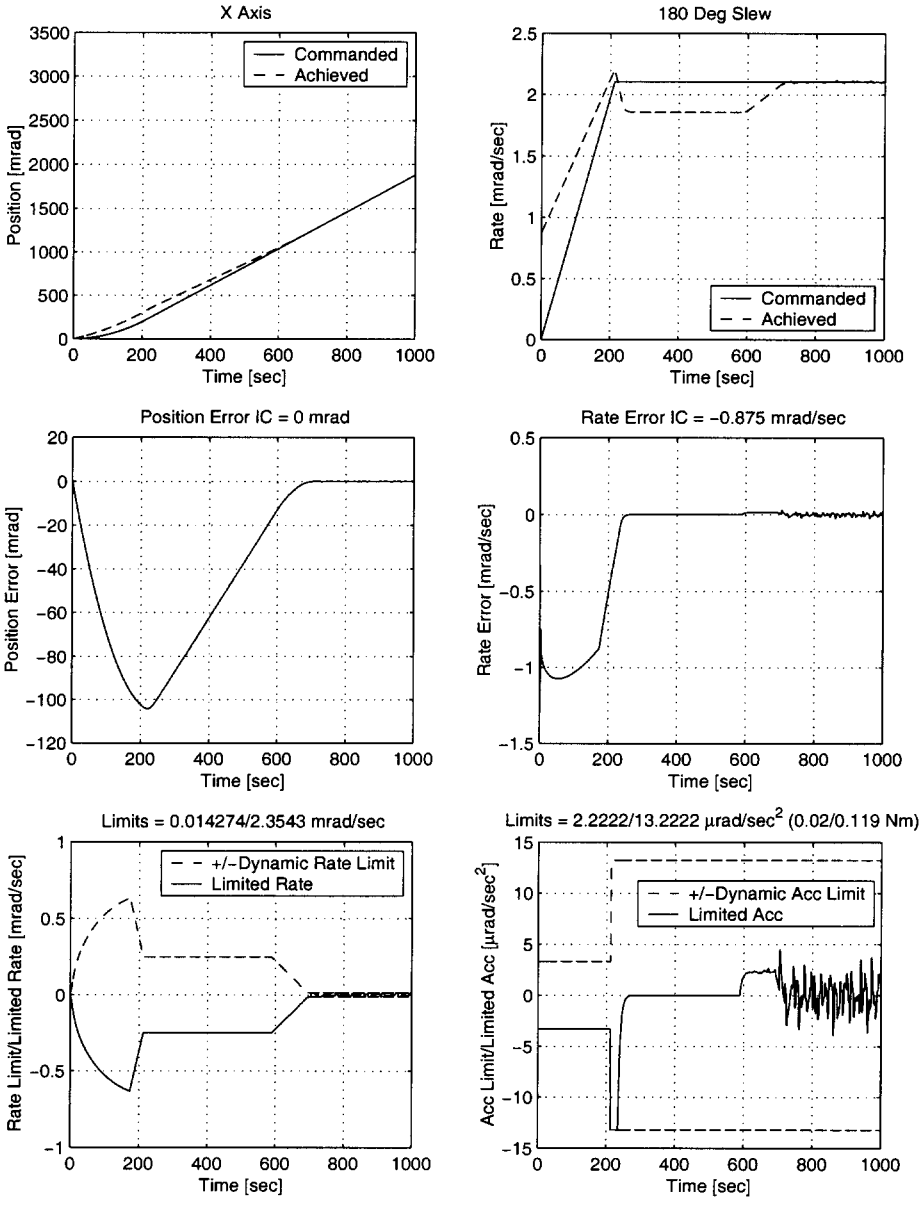

Figure 7 Simulation Results Demonstrating the Dynamic Limiter Features

PERFORMANCE RESULTS

With the design and analysis of RWAC now complete, pointing performance predictions from simulation results are presented as well as some initial results gleaned from in-flight telemetry.

A simulation of the full 3-axis RWAC controlling a flexible dynamics model of the S/C including propellant slosh was used to predict pointing performance. Figure 8 summarizes some of the simulation predictions from a set of 5 Monte Carlo runs. The composite curves are the Root Mean Square of the results of the individual runs (slews) and they indicate that the S/C attitude motion is quieter than required, even when removing margin allocated to S/C instrument motion outside the bandwidth of the controller.

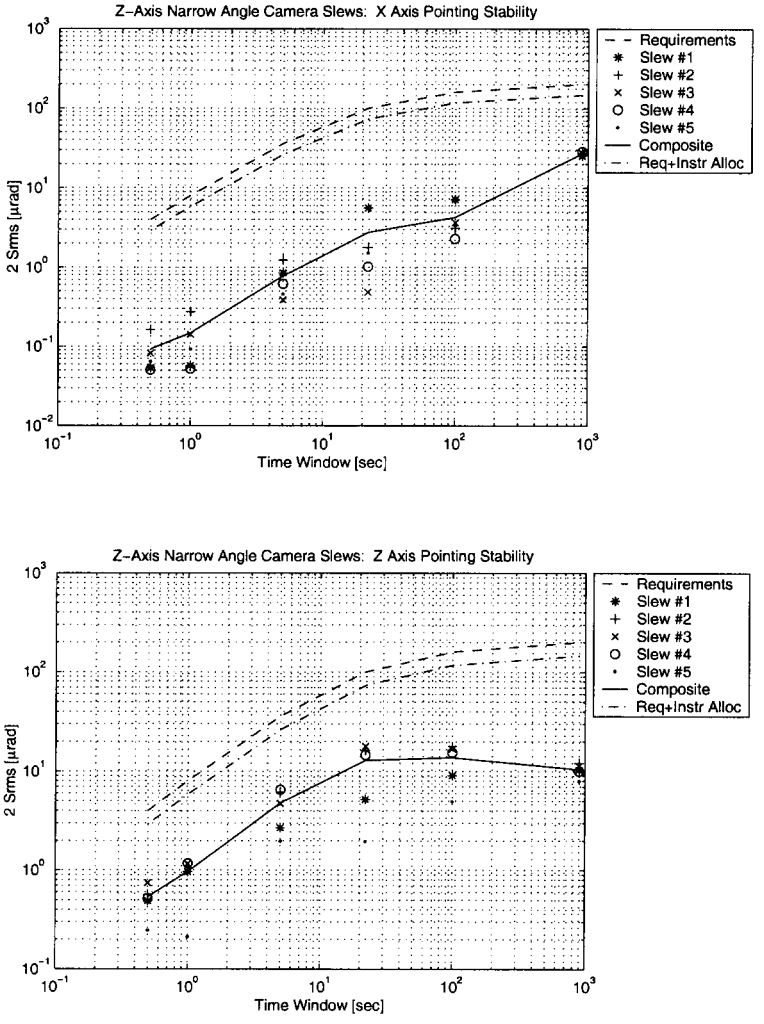

Figure 8 Pointing Stability Simulation Results

Cassini was launched in 1997 and has spent the majority of its flight time using the RCS for attitude control. The original plan was to use the RCS during the cruise to Saturn and then to use the RWA for precision pointing during the orbital tour at Saturn. However, the RWA has now also been used to control attitude in flight, accumulating about a year of flight time. As such, some telemetry has been analyzed to determine RWAC pointing performance in flight. Table 2 compares the results of the analysis against requirements. Where available, it is seen that in-flight operation of the RWAC is easily meeting requirements and appears to be performing substantially better than required for this set of data.

Table 2
S/C RMS Per-Axis Pointing Stability[1] Requirements & Preliminary In-Flight Capability[5]

Time Window (sec)	2 σ per axis requirement (μrad)	2 σ per axis capability (μrad)
0.5	4	unavailable
1	8	unavailable
5	36	10
22	100	26
100	160	51
900	200	55
1200	220	56
1 hr	280	56
4 hr	800	unavailable

CONCLUSIONS

This paper has presented the architecture and methodology used for the design of the Cassini spacecraft's Reaction Wheel Attitude Controller. Simulation results were presented that predict that pointing performance will satisfy pointing requirements. Preliminary in-flight pointing performance analysis verifies that the RWAC is easily meeting pointing requirements and seems to be capable of performing substantially better than required.

ACKNOWLEDGMENTS

The research described in this paper was carried out by the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautic and Space Administration.

REFERENCES

1. S. Sirlin and M. San Martin, "A New Definition of Pointing Stability," JPL EM 343-1189, March 6, 1990.

2. R. Rasmussen, G. Singh, D. B. Rathbun, and G. A. Macala, "Behavioral Model Pointing on Cassini Using Target Vectors", presented at the 18th Annual Rocky Mountain Guidance and Control Conference sponsored by the American Astronautical Society, February 1995 and published in volume 88, *Advances in Astronautical Sciences*, 1995, a publication of the American Astronautical Society (re-published by SPIE with permission in 1996).

3. G. Singh, G. Macala, E. Wong, and R. Rasmussen, "A Constraint Monitor Algorithm for the Cassini Spacecraft," *Proceedings of the AIAA, Guidance, Navigation, and Control Conference*, New Orleans, Louisiana, Aug. 11-13, 1997.

4. P.E. Gill, W. Murray, M.A. Saunders, and M.H. Wright, "User's guide for NPSOL (version 4.0): A FORTRAN package for nonlinear programming", Technical report, Systems Optimization Laboratory, Stanford University, 1986.

5. A. Lee and G. Hanover, "Cassini Pointing Performance", Presentation to Cassini Project Science Group, October 9, 2001.

Design of the Reaction Wheel Attitude Control System for the Cassini Spacecraft

Glenn A. Macala

Jet Propulsion Laboratory
California Institute of Technology
M.S. 198-326
4800 Oak Grove Drive
Pasadena, CA 91109

Cassini Spacecraft

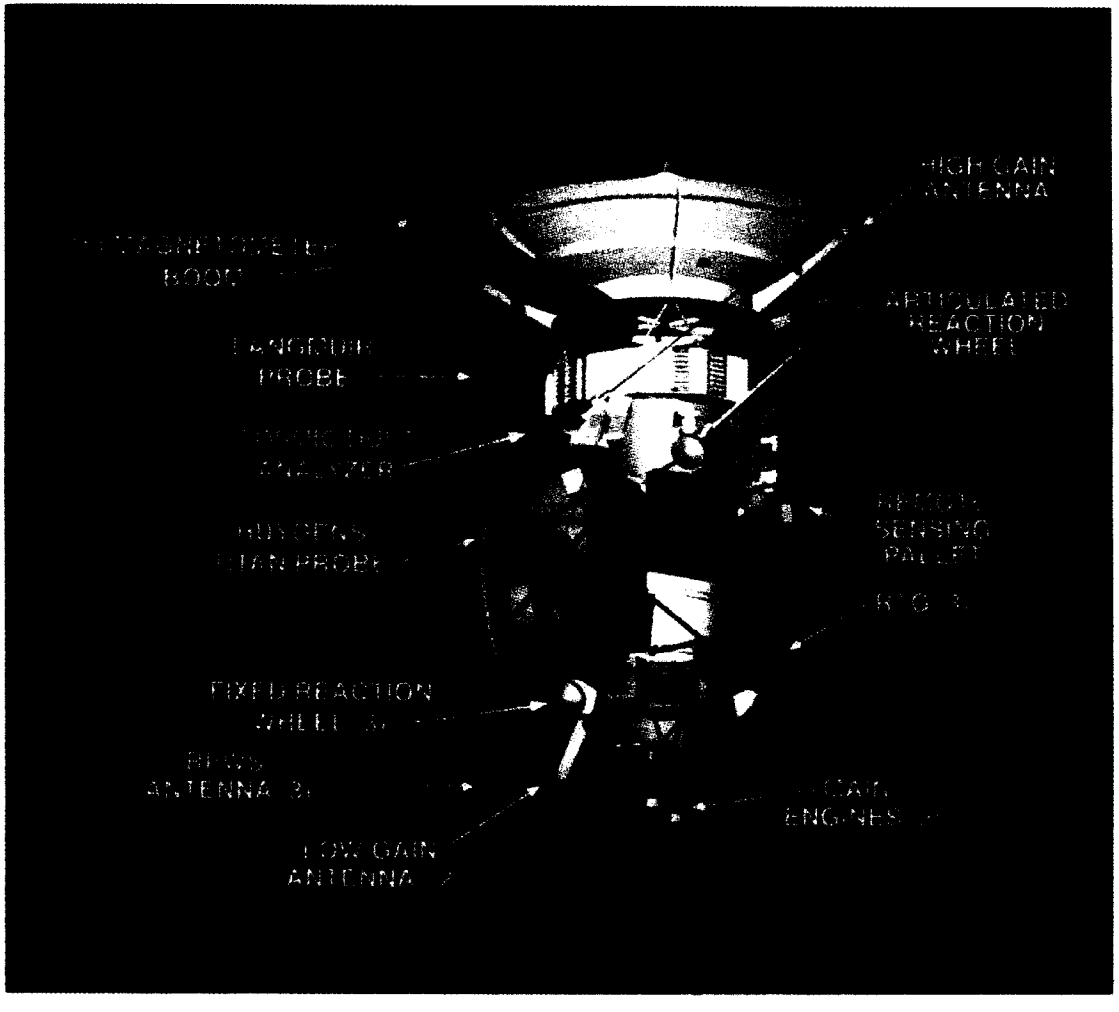

Reaction Wheel Attitude Controller Pointing Requirements

- Relative pointing errors: < 40 μrad per axis at slew rates < 0.13°/sec (X/Y), 0.26°/sec (Z)
- RMS Per-Axis Pointing Stability while Staring

Time Window (sec)	2 σ per axis (μrad)
0.5	4
1	8
5	36
22	100
100	160
900	200
1200	220
1 hr	280
4 hr	800

Reaction Wheel Attitude Controller Block Diagram

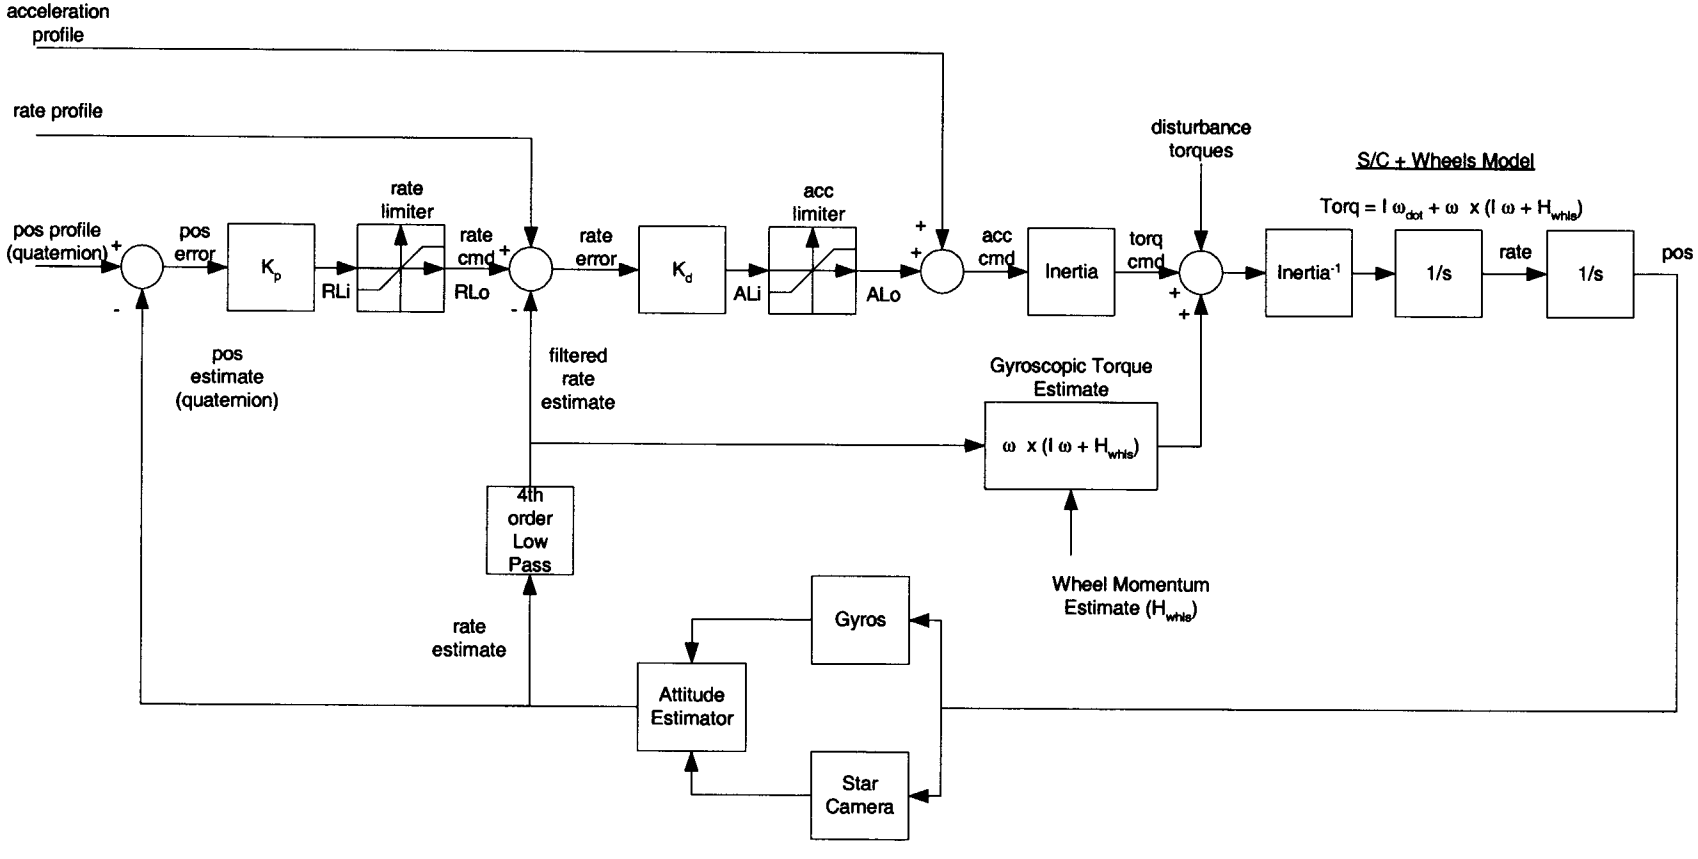

Optimization Using Open Loop Frequency Responses

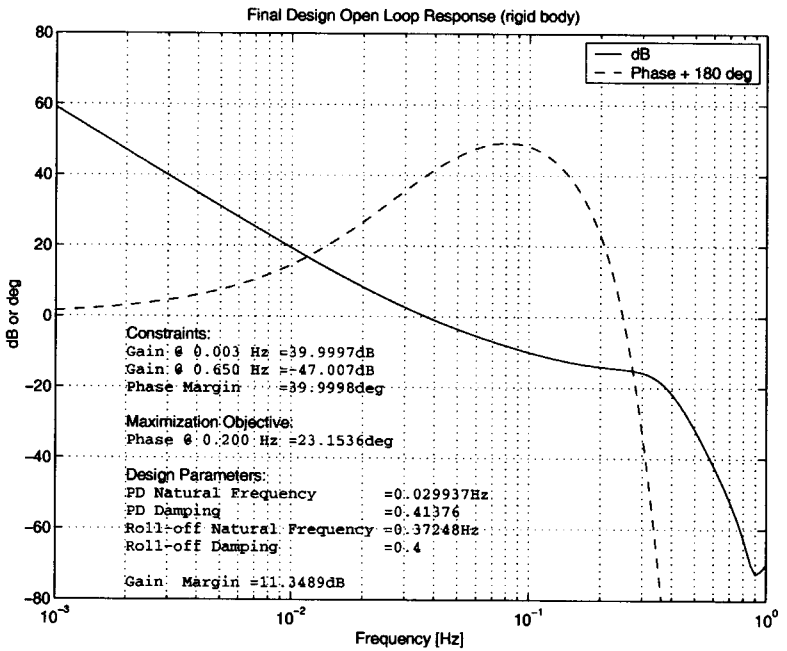

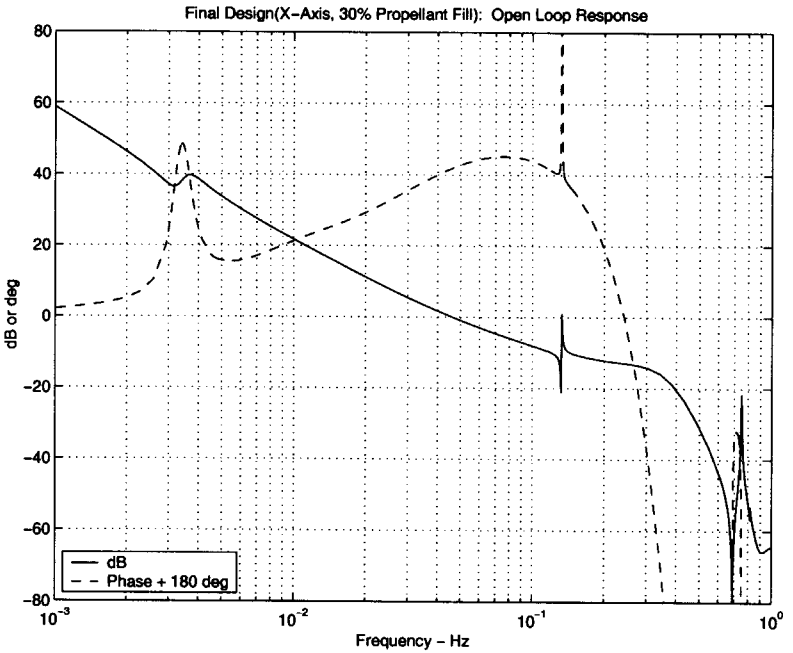

Dynamic Limiters

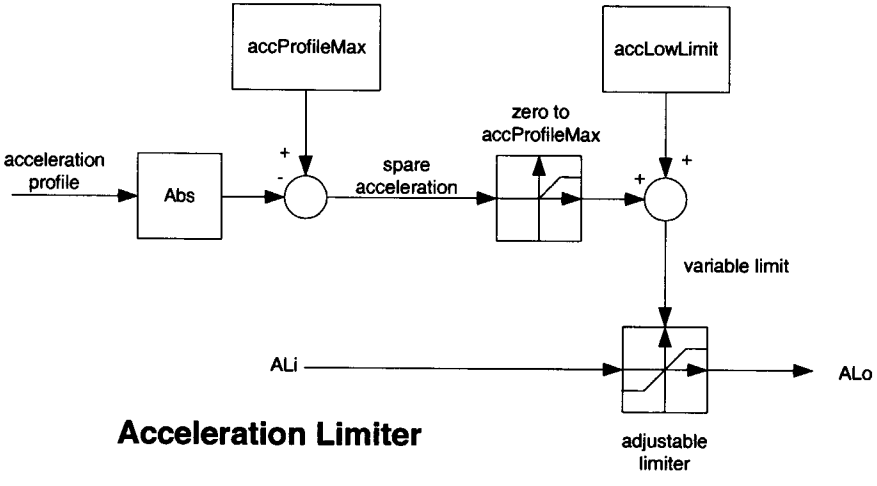

Acceleration Limiter

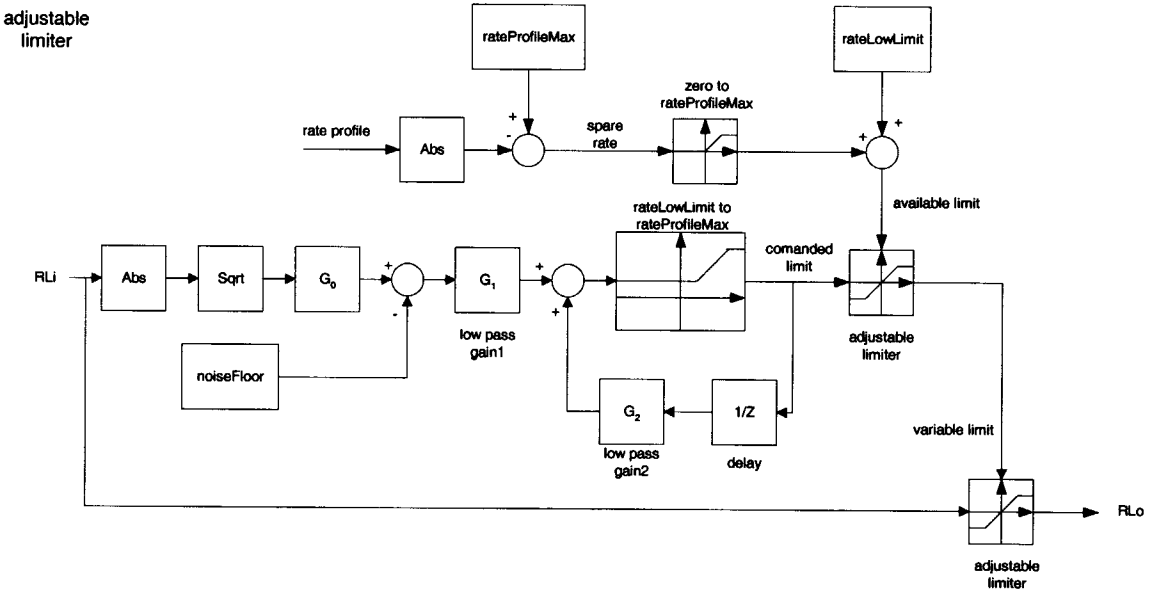

Rate Limiter

Simulation Results Demonstrating Dynamic Limiters

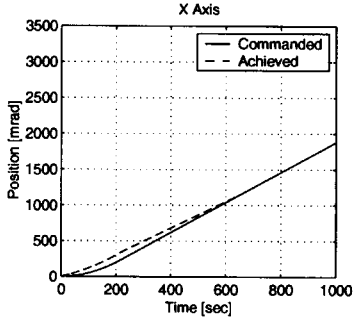

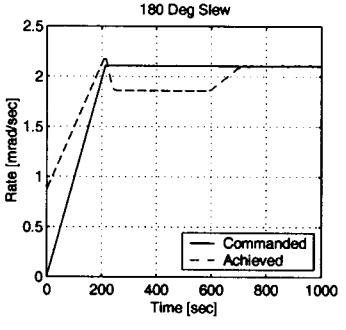

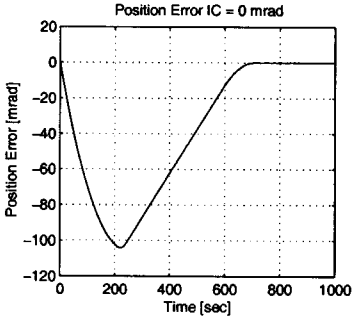

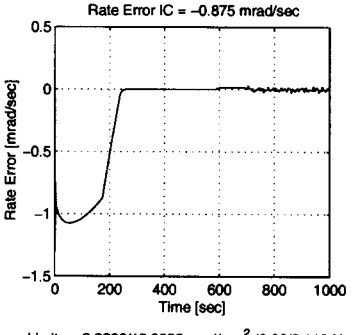

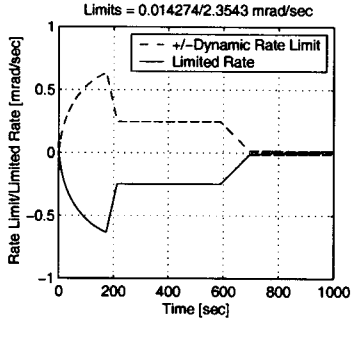

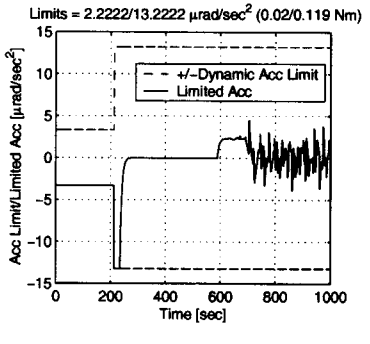

Pointing Stability Simulation Results

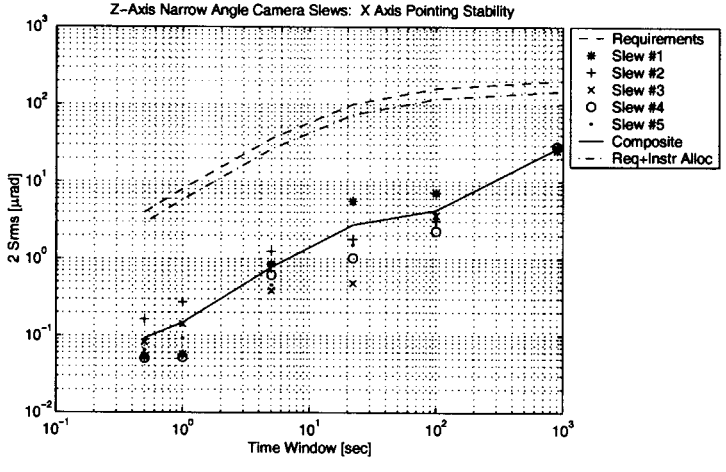

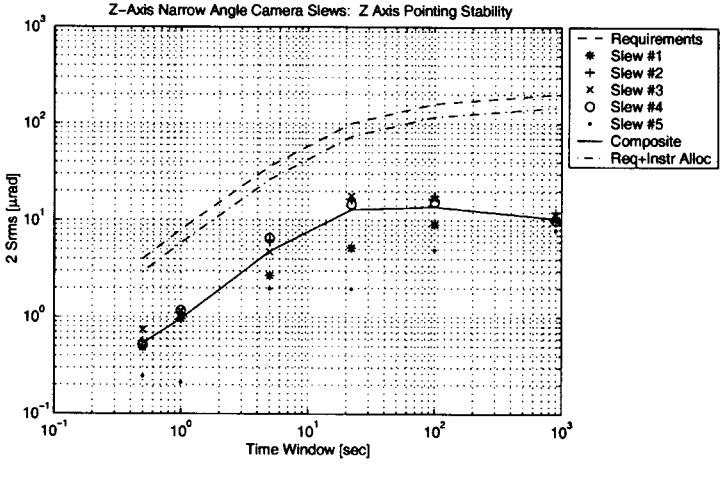

Preliminary Assessment of In-Flight Pointing Stability

Time Window (sec)	2-σ per axis requirement (μrad)	2-σ per axis capability (μrad)
0.5	4	unavailable
1	8	unavailable
5	36	10
22	100	26
100	160	51
900	200	55
1200	220	56
1 hr	280	56
4 hr	800	unavailable

Conclusions

- Presented the architecture and methodology used for the design of the Cassini spacecraft's Reaction Wheel Attitude Controller
- Simulation results predict pointing performance will satisfy pointing requirements
- Preliminary in-flight pointing performance analysis verifies that pointing requirements are being met and seem to be substantially better than requirements in some cases

December 2000: Jupiter Swing-by

"The spacecraft is steadier than any spacecraft I've ever seen," said Dr. Carolyn Porco of the University of Arizona, team leader for the camera on Cassini. "It's so steady, the images are unexpectedly sharp and clear, even in the longest exposures taken and most challenging spectral regions."